



INVESTOR IN PEOPLE

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.


Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.



Signed

Dated

3 November 2003



The
Patent
Office

1/77

Patents Act 1977
Rule 16

01SEP03 E833950-1 000611
P01/7700 0.00-0320386.6

Request for grant of a patent

30 AUG 2003

The Patent Office
Concept House
Cardiff Road
Newport
South Wales NP10 8QQ

1.	Your reference	GB920030077GB1		
2.	Patent application number (The Patent Office will fill in this part)	0320386.6		
3.	Full name, address and postcode of the or of each applicant (underline all surnames)	INTERNATIONAL BUSINESS MACHINES CORPORATION Armonk New York 10504 United States of America		
	Patents ADP number (if you know it)	00519637001		
	If the applicant is a corporate body, give the country/state of its incorporation	State of New York United States of America		
4.	Title of the invention	A METHOD, APPARATUS AND COMPUTER PROGRAM FOR EXECUTING A PROGRAM		
5.	Name of your agent (if you have one)	J D Williams		
	"Address for Service" in the United Kingdom to which all correspondence should be sent (including the postcode)	IBM United Kingdom Limited Intellectual Property Department Hursley Park Winchester Hampshire SO21 2JN		
	Patents ADP number (if you know it).	04474847002 919001		
6.	If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (if you know it) the or each application number	Country	Priority App No (if you know it)	Date of filing (day/month/year)
7.	If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date or the earlier application	No of earlier application	Date of filing (day/month/year)	

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer 'Yes' if:
a) any applicant named in part 3 is not an inventor, or
b) there is an inventor who is not named as an applicant, or
c) any named applicant is a corporate body.) Yes

9. Enter the number of sheets for any of the following items you are filing with this form. Do not count copies of the same document

Continuation sheets of this form

Description 8

Claim(s) 2

Abstract 1

Drawing(s) 4 *24*

10. If you are also filing any of the following, state how many against each item.

Priority documents

Translations of priority documents

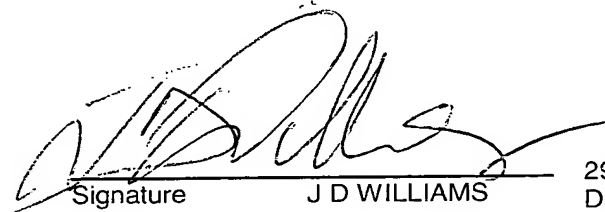
Statement of inventorship and right to grant of a patent (Patents Form 7/77)

Request for preliminary examination and search (Patents Form 9/77)

Request for substantive examination (Patents Form 10/77)

Any other documents (please specify)

11. I/We request the grant of a patent on the basis of this application




Signature

J D WILLIAMS

29 August 2003
Date

12. Name and daytime telephone number of person to contact in the United Kingdom J N Watson
01962 818955




7/77

Patents Act 1977
Rule 15

30 AUG 2003

Statement of inventorship and of right to grant of a patent

The Patent Office
Concept House
Cardiff Road
Newport
South Wales NP10 8QQ

1.	Your reference	GB920030077GB1
2.	Patent application number (if you know it)	0320386.6
3.	Full name of the or of each applicant	INTERNATIONAL BUSINESS MACHINES CORPORATION
4.	Title of invention	A METHOD, APPARATUS AND COMPUTER PROGRAM FOR EXECUTING A PROGRAM
5.	State how the applicant(s) derived the right from the inventor(s) to be granted a patent	By employment and by agreement
6.	How many, if any, additional Patents Forms 7/77 are attached to this form?	
7.	I/We believe that the person(s) named over the page (and on any extra copies of this form) is/are the inventor(s) of the invention which the above patent application relates to.	
	 <div style="display: flex; justify-content: space-between;"> Signature J D WILLIAMS </div> <div style="text-align: right;"> 29 August 2003 Date </div>	
8.	Name and daytime telephone number of person to contact in the United Kingdom	J N Watson Tel: 01962 818955

Patents Form 7/77

Enter the full names, addresses and postcodes of the inventors in the boxes and underline the surnames

Adam Dean ILEY
(UK resident)
c/o IBM United Kingdom Limited
Intellectual Property Law
Hursley Park
Winchester
Hampshire SO21 2JN
England

Patents ADP number (if known)

08592099001

John SCOTT
(UK resident)
c/o IBM United Kingdom Limited
Intellectual Property Law
Hursley Park
Winchester
Hampshire SO21 2JN
England

Patents ADP number (if known)

08704249001

If there are more than three inventors, please write their names and addresses on the back of another Patents Form 7/77 and attach it to this form

REMINDER

Have you signed the form?

Patents ADP number (if known)

A METHOD, APPARATUS AND COMPUTER PROGRAM FOR EXECUTING A PROGRAM

FIELD OF THE INVENTION

5 The invention relates to speculative pre-execution of portions of a computer program.

BACKGROUND OF THE INVENTION

10 Computers have proliferated into all aspects of society and in today's increasingly competitive market-place, the performance of not only the machines themselves but also the software that runs on these machines, is of the utmost importance. Software developers are therefore continually looking for methods to improve the execution efficiency of the code
15 (programs) they produce in order to meet the high expectations of software users.

 One such method is by inserting pre-execution instructions into source code such that execution of such instructions cause a portion of the
20 program defined by the source code to be pre-executed. This is described in US Patent Application Publication US 2002/0055964.

 Further, US Patent Application Publication US 2002/0144083 describes a processor using spare hardware contexts to spawn speculative threads such
25 that data is pre-fetched in advance of a main thread.

 Another known method is "branch prediction" (also mentioned in US 2002/0055964). Within a program there are typically a number of branch points. These are points which can return one of a finite number of
30 results. Prediction techniques are used to determine the likely return result such that a branch point's subsequent instructions can be pre-executed on this assumption. "if...else" statements and "case" statements are two well known examples of branch points.

35 There are a number of branch prediction techniques known in the industry. Such techniques are common in RISC and processor architectures (e.g. The pSeries architecture).
See
also [www.mtl.t.u-tokyo.ac.jp/~niko/Downloads/chitaka-EuroPar2001-PerThreadPr
40 edictor.pdf](http://www.mtl.t.u-tokyo.ac.jp/~niko/Downloads/chitaka-EuroPar2001-PerThreadPredictor.pdf) which presents a hardware scheme for improving branch prediction accuracy.

 Software schemes also exist. A paper "Static Correlated Branch Prediction" by Cliff Young and Michael D Smith (ACM Transactions on

Programming Languages and Systems, Vol. 21. No ?, ??? 1999, Pages 111-159) describes how the repetitive behaviour in the trace of all conditional branches executed by a program can be exploited by a compiler. Another paper "A Comparative Analysis of Schemes for Correlated Branch Prediction" by Cliff Young, Michael D Smith and Nicholas Gloy (published in the Proceedings of the 22nd Annual International Symposium on Computer Architecture, June 1995) presents a framework that categorizes branch prediction schemes by the way in which they partition dynamic branches and by the kind of predictor they use.

The paper "Understanding Backward Slices of Performance Degrading Instructions" by C Zilles and G Sohi (published in the proceedings of the 27th Annual International Symposium on Computer Architecture (ISCA - 2000), June 12-14 2000) discusses the small fraction of static instructions whose behaviour cannot be anticipated using current branch predictors and caches. The paper analyses the dynamic instruction stream leading up to these performance degrading instructions to identify the operations necessary to execute them early.

Another paper "The Predictability of Computations that Produce Unpredictable Outcomes" by T Aamodt, A Moshovos and P Chow (an update of the paper that appeared in the Proceedings of the 5th Workshop on Multithreaded Execution, Architecture, and Compilation - pages 23-34, Austin, TX, December 2001) studies the dynamic stream of slice traces that foil existing branch predictors and measures whether these slices exhibit repetition.

"Speculative Data-Driven Multithreading" by Amir Roth and Gurindar Sohi (appearing in the Proceedings of the 7th International Conference on High Performance Computer Architecture (HPCA-7), Jan 22-24, 2001) describes the use of speculative data-driven multithreading (DDMT) for coping with mispredicted branches and loads that miss in the cache.

It is also known for the programmer to be able to provide branch prediction pragma - see <http://www.geocrawler.com/archives/3/357/1993/7/0/1992785/>.

Whilst branch prediction techniques are known, there is however a need in the industry for more efficient processing of software functions as opposed to branch points.

SUMMARY

Accordingly the invention provides a method for executing a program comprising a function call and one or more subsequent instructions, the method comprising the steps of: processing, on a first thread, a function defined by the function call, the function having one or more programmer predefined typical return values; for each predefined return value, pre-processing, on an additional thread, the one or more subsequent instructions assuming that the function returned that pre-defined return value, thereby enabling said processor, on completion of processing said function, to make use of the pre-processing completed by the additional thread which used the actual return value.

Thus the present invention enables a programmer to define typical return values for a function such that the function can be pre-processed ahead of a main thread. Assuming that the function does actually return one of the predefined return values, performance can be much improved.

Note, preferably the additional threads operate in parallel.

Preferably the program comprises a plurality of subsequent instructions defining one or more additional functions and the plurality of subsequent instructions are pre-processed on each additional thread until a function is reached which is of external effect. Once such a function is reached by an additional thread that thread preferably blocks (waits) on said function until the actual return value is determined by the first thread.

Preferably each additional thread also blocks on reaching a function which is affected by an external event.

According to one aspect the invention provides an apparatus for executing a program comprising a function call and one or more subsequent instructions, the apparatus comprising: means for processing, on a first thread, a function defined by the function call, the function having one or more programmer predefined typical return values; means for pre-processing for each predefined return value, on an additional thread, the one or more subsequent instructions assuming that the function returned that pre-defined return value, thereby enabling said processor, on completion of processing said function, to make use of the pre-processing completed by the additional thread which used the actual return value.

The invention may be implemented in computer software.

According to another aspect, the invention provides a compiler for generating a computer program comprising a function call defining a function, having one or more programmer predefined typical return values, and one or more subsequent instructions, the compiler comprising means for generating executable code, said executable code for instructing a computer to process on a first thread the function and to pre-process, for each defined typical return value, on an additional thread the one or more subsequent instructions assuming that the function returned that pre-defined return value, thereby enabling said processor, on completion of processing said function, to make use of the pre-processing completed by the additional thread which used the actual return value.

It will be appreciated that the term compiler is intended to cover the whole compilation process optionally including linking.

BRIEF DESCRIPTION OF THE DRAWINGS

A preferred embodiment of the present invention will now be described, by way of example only, and with reference to the following drawings:

Figure 1 illustrates an extract of psuedo code incorporating the new construct provided by a preferred embodiment of the present invention;

Figure 2a shows the processing of spawned pre-execution threads in accordance with a preferred embodiment of the present invention;

Figure 2b shows the processing of a main thread in accordance with a preferred embodiment of the present invention; and

Figure 3 illustrates the operation of a compiler in accordance with a preferred embodiment of the present invention.

DETAILED DESCRIPTION

It has been observed that within a program certain tasks (functions) require substantial amounts of processing time but frequently return the same result. In order to exploit this observation a new construct is preferably incorporated into existing programming languages. This construct enables programmers to mark certain functions as "restricted". In this context, the keyword "restricted" preferably means that the marked function does not effect the global environment (e.g. by outputting to a file) and the syntax associated with the new keyword permits the values most commonly returned by the function to be specified by the programmer as

part of the function's signature. Further preferably, a "restricted" function is not itself affected by the global environment. In other words, it always operates in the same way regardless of the results produced by other "restricted" functions.

Figure 1 shows an extract of pseudo code from a library program incorporating the new "restricted" keyword in accordance with a preferred embodiment of the present invention. The extract of library program shown includes two main functions: `overdue`; and `send_letter_to_printer`. The `overdue` function is marked as "restricted" since it does not affect the global environment. By contrast the `send_letter_to_printer` function results in printer output and does not therefore have the "restricted" keyword associated with it.

From the code extract, it can be seen that the `overdue` function checks the status of each user's book to determine whether that book is: not yet due back at the library; is late back; or is very late back. If a user's book is not overdue, then the function does no processing in relation to that user. On the other hand, if a user's book is either late or very late, then the `remind_late` or `remind_very_late` function is called as appropriate.

Whilst the `overdue` function itself is thus relatively fast, both `remind` functions have long and complicated processing to do on behalf of the user in relation to which that function is called. This processing involves looking up the user's address; the name of the overdue book; the number of days the book is overdue by; and the list of those currently waiting for the book. If the book is very late, then the user's borrower history must also be checked. Further, in both cases the outstanding fine has to be calculated and the appropriate letter text retrieved. All this information is then used to build an appropriate letter in memory for eventual dispatch to the user.

Whilst the processing of both `remind` functions is long and complicated, this processing also does not affect the global environment. Values are retrieved and held in volatile memory, but no data is inserted, updated, deleted or output to non-volatile memory, an external device etc.. Thus these functions can also be marked as "restricted", although in this instance it is not appropriate to associate either function with typical return values.

Once letters have been built in non-volatile memory for all user's with overdue books, then these letters are sent to the printer via the

"send_to_printer" function. This function is not marked as "restricted" since it does effect the global environment.

The execution of code including the new "restricted" keyword will now be described with reference to figures 2a and 2b.

Figure 2a shows the processing of pre-execution threads in accordance with a preferred embodiment of the present invention. Upon encountering a restricted function having typical return values defined (as described above), a pre-execution thread is spawned for each such return value (step 100). For each such pre-execution thread, instructions subsequent to the restricted function are executed as if the restricted function did indeed return the value associated with the particular pre-execution thread (step 110). In other words, the restricted function is not actually executed. Instead, for each pre-execution thread, it is assumed that the function returned one of the predefined values. Each pre-execution thread then continues executing instructions until a non-restricted function is encountered (step 120). As discussed above, non-restricted functions affect the global environment via, for example, updating data; inserting data; deleting data; or outputting results. Thus each pre-execution thread then blocks on the non-restricted function until the true result of the original "restricted" function is determined by a main thread (step 130).

Note, as alluded to with reference to figure 1, not all "restricted" functions have typical return values associated therewith. For example, the remind functions do not since they rely upon the results returned by the overdue function.

Further, rather than spawning pre-execution threads, a thread pool may be used.

Figure 2b shows the processing of a main thread in accordance with a preferred embodiment of the present invention. The main thread processes a "restricted" function having typical return values defined (step 200). Upon determining the result actually returned by this function, the main thread determines whether this result corresponds to one of the defined return values associated with the "restricted" function (step 210). Assuming that the return value does correspond to one of the defined return values, then the main thread is terminated and execution skips to the non-restricted function (step 220). Execution then continues using the pre-execution thread associated with the actual return value (step 230). All other pre-execution threads are terminated (step 240).

Thus by enabling the programmer to define functions with non-global effect/as not affected by the global environment and also typical return values for such functions, it is possible to speculatively pre-execute code. Assuming that the speculation proves correct, program execution performance can be dramatically improved - a pre-execution thread will have preferably performed the long and complicated processing in the background whilst the main thread is performing other tasks.

Note, in one embodiment the main thread is not finally terminated until it is verified that an appropriate pre-execution does exist. Indeed it may be the main thread that is responsible for terminating those pre-execution threads that are not associated with the correct return value.

Another example of a system in which the invention should prove useful is a menu system in which a program will display a number of menu options and then wait for the user to choose one. In accordance with the "restricted" construct defined by a preferred embodiment of the present invention, the programmer can define the options most likely to be selected and then the program can pre-execute each of those options as far as it can (i.e. until a global function is encountered).

As discussed above, the functionality of the present invention is preferably achieved by modification of existing programming languages. Executable programs are typically produced from compiled source code. The compilation process is thus modified such that the meaning of "restricted" keyword is understood and such that appropriate executable code is generated as a result of the compilation process.

Thus for completeness the operation of a compiler in accordance with a preferred embodiment of the present invention is described with reference to figure 3.

A compiler 310 is provided with a program's source code 300 as input. The compiler processes this source code to produce object code 320 and this is then passed to a linker 330 which uses this code 320 to produce an executable 340.

Typically, there are three stages to the compilation process: lexical analysis; syntax analysis; and code generation. During the lexical analysis, symbols (e.g. alphabetic characters) are grouped together to form tokens. For example the characters P R I N T are grouped to form the command (token) PRINT. In some systems, certain keywords are replaced by

shorter, more efficient tokens. This part of the compilation process also verifies that the tokens are valid.

5 In accordance with a preferred embodiment of the present invention, the lexical analyser is therefore modified to recognise "restricted" as a keyword and also to recognise expected return values when the programmer provides them.

10 Next, the syntax analyser checks whether each string of tokens forms a valid sentence. Again the syntax analyser is preferably modified to recognise that "restricted" keyword and the predefined typical return values are valid.

15 Finally, the code generation stage produces the appropriate object code. The code generator is thus also preferably modified to recognise the new "restricted" construct such that the appropriate object code is generated for any program employing the new construct (i.e. to achieve the result discussed with reference to figures 2a and 2b.)

20 It is assumed that a person skilled in the art of compiler development will be familiar with the above process and thus this will not be discussed in any further detail.

CLAIMS

1. A method for executing a program comprising a function call and one or more subsequent instructions, the method comprising the steps of:

processing, on a first thread, a function defined by the function call, the function having one or more programmer predefined typical return values;

for each predefined return value, pre-processing, on an additional thread, the one or more subsequent instructions assuming that the function returned that pre-defined return value,

thereby enabling said processor, on completion of processing said function, to make use of the pre-processing completed by the additional thread which used the actual return value.

2. The method of claim 1, wherein the program comprises a plurality of subsequent instructions defining one or more additional functions, the method further comprising:

pre-processing on each additional thread the plurality of subsequent instructions until a function is reached which is of external effect; and

blocking on said function having external effect until the actual return value is determined by the first thread.

3. The method of claim 2, wherein the blocking step also blocks on reaching a function which is affected by an external event.

4. An apparatus for executing a program comprising a function call and one or more subsequent instructions, the apparatus comprising:

means for processing, on a first thread, a function defined by the function call, the function having one or more predefined typical return values;

means for pre-processing for each predefined return value, on an additional thread, the one or more subsequent instructions assuming that the function returned that pre-defined return value,

thereby enabling said processor, on completion of processing said function, to make use of the pre-processing completed by the additional thread which used the actual return value.

5. The apparatus of claim 4, wherein the program comprises a plurality of subsequent instructions defining one or more additional functions, the apparatus further comprising:

5 means for pre-processing on each additional thread the plurality of subsequent instructions until a function is reached which is of external effect; and

10 means for blocking on said function having external effect until the actual return value is determined by the first thread.

6. The apparatus of claim 5, wherein the blocking means is operable to also block on reaching a function which is affected by an external event.

15 7. A computer program comprising program code means adapted to perform, when said program is run on a computer, the method of any of claims 1 to 3.


20 8. A compiler for generating a computer program comprising a function call defining a function, having one or more programmer predefined typical return values, and one or more subsequent instructions, the compiler comprising means for generating executable code, said executable code for instructing a computer to process on a first thread the function and to pre-process, for each defined typical return value, on an additional thread the one or more subsequent instructions assuming that the function returned
25 that pre-defined return value, thereby enabling said processor, on completion of processing said function, to make use of the pre-processing completed by the additional thread which used the actual return value.

ABSTRACT**A METHOD, APPARATUS AND COMPUTER PROGRAM FOR EXECUTING A PROGRAM**

5 There is provided a method for executing a program comprising a
function call and one or more subsequent instructions. The method
comprises processing, on a first thread, a function defined by the function
call, the function having one or more programmer predefined typical return
values. For each predefined return value, the one or more subsequent
10 instructions are pre-processed on an additional thread assuming that the
function returned that pre-defined return value. In this way the
processor, on completion of processing said function is able to make use of
the pre-processing completed by the additional thread which used the actual
return value.

15

THIS PAGE BLANK (USPTO)


restricted (late, very late) library()

{

while (more users)

{

determine status of user's book

if not overdue

{

do nothing

}

if late

{

remind_late_user(user)

}

if very late

{

remind_very_late(user)

}

}

while(more letters)

{

send_letter_to_printer(letter)

}

}

restricted remind_late_user(user)

{

 lookup user's address, book name; number of days overdue by; outstanding
 fine; waiting list; borrower history
 calculating fine due
 retrieve late text
 build letter in volatile memory

}

restricted remind_very_late_user(user)

{

 lookup user's address; book name; number of days overdue by, waiting list
 calculating fine due
 retrieve very late text
 build letter in volatile memory

}

send_letter_to_printer(letter)

{

print letter}

THIS PAGE BLANK (USPTO)

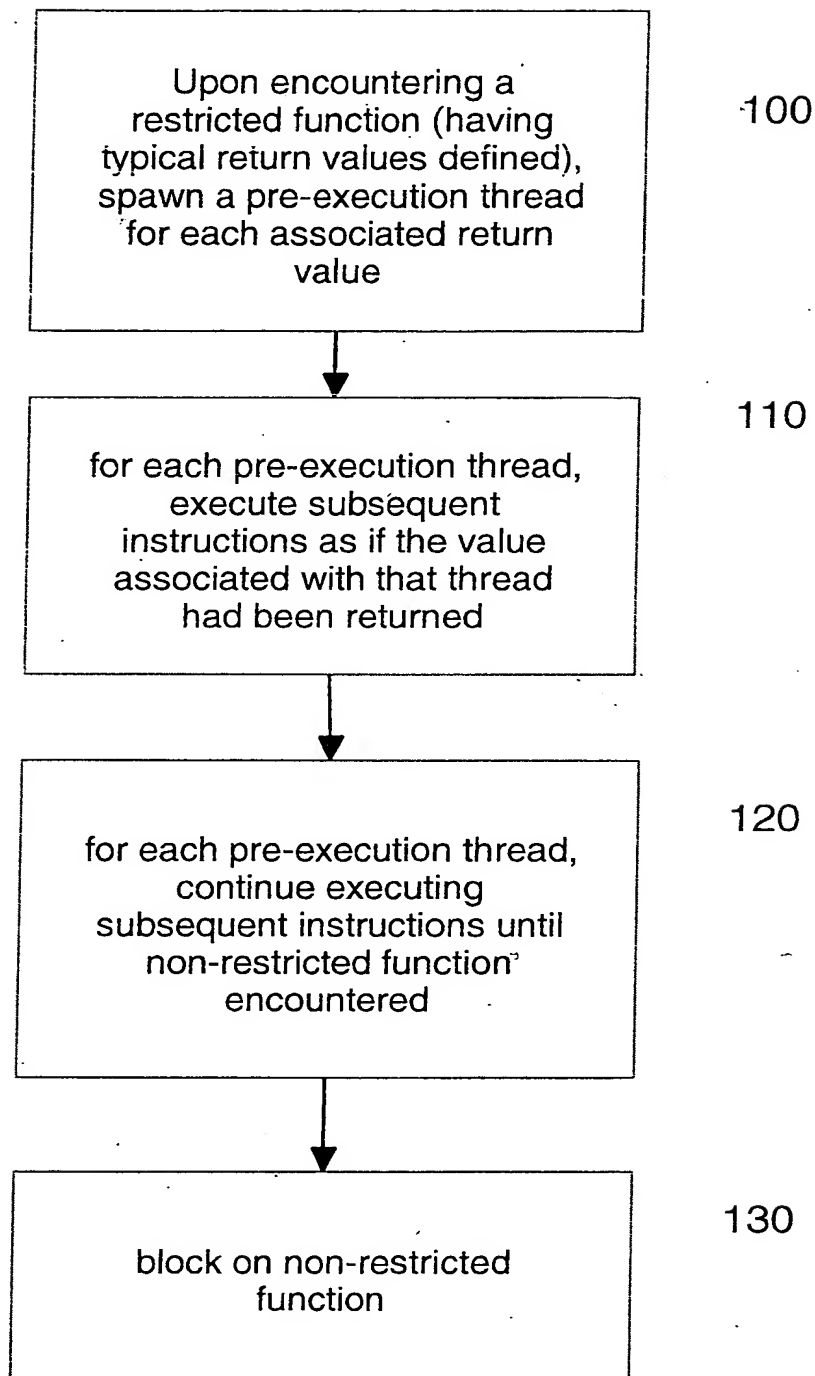


Figure 2a

THIS PAGE BLANK (USPTO)

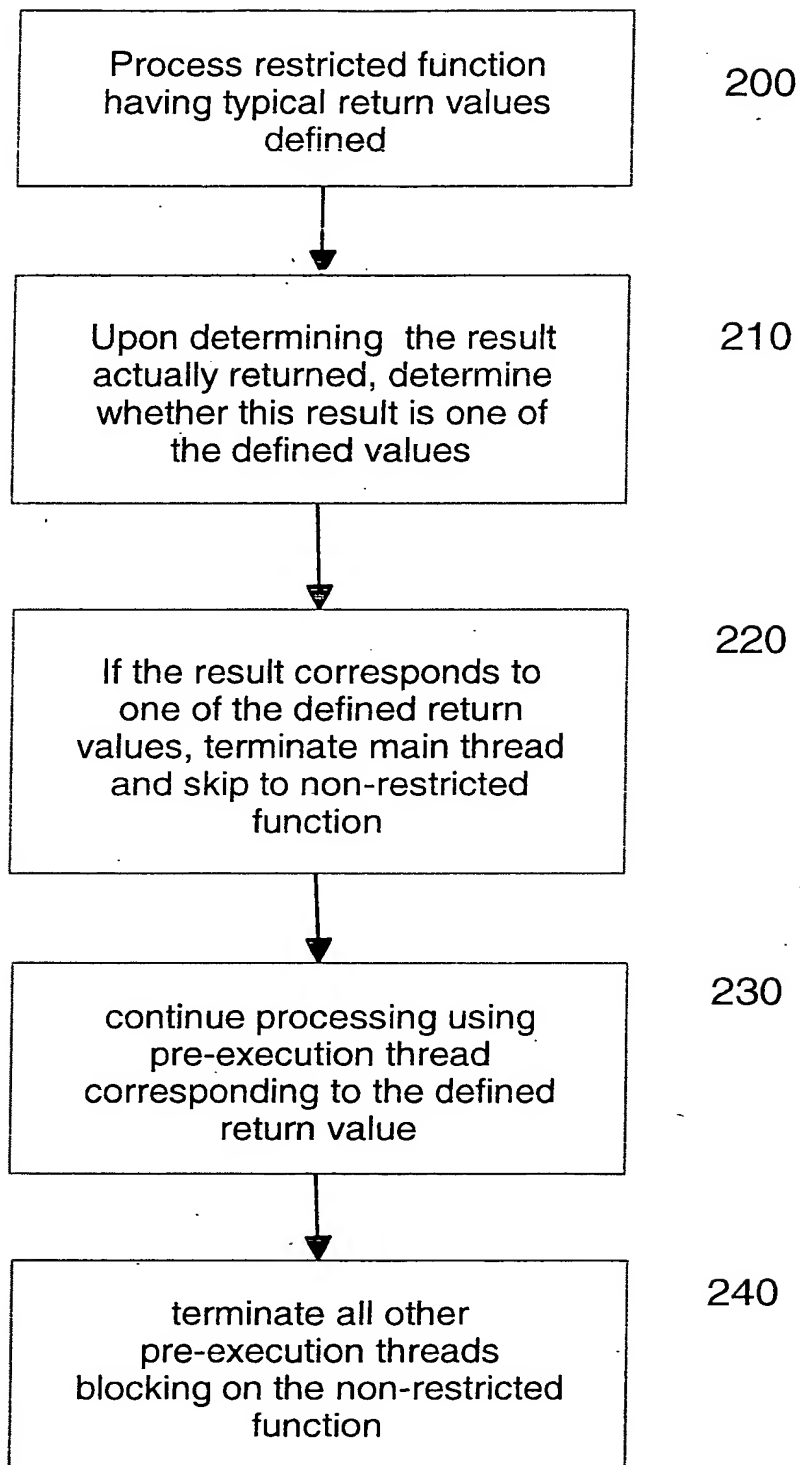


Figure 2b

THIS PAGE BLANK (USPTO)




Figure 3

THIS PAGE BLANK (USPTO)